



ELECTRONICS RESEARCH LABORATORY

Information Technology Division

RESEARCH NOTE
ERL-0553-RN

TECHNIQUES FOR KNOWLEDGE ACQUISITION

by

S. J. Cosgrove

SUMMARY

Knowledge acquisition is commonly regarded as the major bottle-neck in Expert System (ES) development. Numerous techniques and commercial packages have been built in order to widen this bottle-neck. Unfortunately, these have had little success, except in simple, highly confined domains, therefore knowledge acquisition is still regarded as a "black art". This paper critically analyses current acquisition techniques and suggests future research directions.

© COMMONWEALTH OF AUSTRALIA 1991

June 91

COPY NO.

APPROVED FOR PUBLIC RELEASE

POSTAL ADDRESS: Director, Electronics Research Laboratory, PO Box 1600, Salisbury, South Australia, 5108.

ERL-0553-RN

This work is Copyright. Apart from any use as permitted under the Copyright Act 1968, no part may be reproduced by any process without written permission from the Australian Government Publishing Service. Requests and enquiries concerning reproduction and rights should be directed to the Manager, AGPS Press, GPO Box 84, Canberra ACT 2601.

CONTENTS

1	INTRODUCTION	1
2	EXPERT SYSTEM KNOWLEDGE-BASE FORMALISMS.....	1
2.1	Structural Knowledge.....	1
2.2	Empirical Associations.....	1
2.3	Causal Models.....	2
2.4	Strategic Knowledge	2
3	KNOWLEDGE ACQUISITION BY ARTICULATION.....	2
3.1	Concept Sorting	2
3.2	Critical Incident.....	2
3.3	Conversation Theory and Focussed Interviewing.....	3
3.4	Problem Solving	3
3.5	Repertory Grids.....	3
4	KNOWLEDGE ARTICULATION BY INDUCTION.....	4
4.1	The ID3 Algorithm.....	4
4.2	The Disadvantages of the ID3 Algorithm	7
5	STRUCTURED AUTOMATED KNOWLEDGE ACQUISITION.....	8
5.1	MORE.....	8
6	CONCLUSION AND FUTURE DIRECTIONS.....	9
7	REFERENCES	10

FIGURES

1	Complexity Relationships of Rule Induction Sets.....	7
2	Example of a Decision Tree Structure	7

THIS IS A BLANK PAGE

1 Introduction

An ES embodies a human expert's domain-specific knowledge, judgment and experience, the acquisition of which is refined over a number of years and forms an integral part of ES development. Knowledge acquisition techniques have evolved into three major categories, the most commonly used being the highly iterative articulation approach. This involves interaction between the knowledge engineer and an expert (or group of experts), in which a two tier hierarchy (the level of reasoning, and the level of underlying structure and function) is iteratively acquired. The result is a model or schema encompassing what is thought of as relevant knowledge.

In contrast to the articulation approach, automated rule induction algorithms [Quinlan83] have been devised whereby the expert supplies a number of case studies by way of domain specific attributes, values and concluding hypotheses. These case studies are then analysed and a set of rules are induced.

The third approach is shell automated knowledge acquisition. This is basically a systematic top-down question and answer process in which a hierarchical rule-based structure is created.

In section 2, ES knowledge-base formalisms are defined. Section 3 then describes knowledge acquisition by articulation. Rule induction theory is discussed in section 4, after which section 5 outlines automated acquisition techniques. Finally, section 6 indicates future directions based upon the limitations of the discussed approaches.

2 Expert system Knowledge - base formalisms

A knowledge-base is a model of a human's understanding of a highly confined domain of expertise. Knowledge acquisition is a methodology for extracting this expertise, from which a software model can be constructed. The developed model will generally consist of numerous types of knowledge whose function, for example, may be that of: disregarding or supporting hypotheses, emulating the structure of class hierarchies or defining the search strategies applicable to the overall domain of expertise.

Before attempting to acquire domain specific knowledge, it is essential that the knowledge engineer fully understands the types of knowledge that can be uncovered. In addition, he must also be able to identify and construct relational models between the various types of knowledge formalisms that have been discovered. Knowledge-based formalisms generally comprise a four tier hierarchy, the components of which are commonly known as: Structural Knowledge, Empirical Associations, Causal Models and Strategic Knowledge.

2.1 Structural Knowledge

Structural Knowledge is the lowest level of the knowledge classification hierarchy and consists of concepts easily obtainable from manuals and books. These concepts may include classification information, basic interconnections and interrelations between a domain's subclasses [Johnson83].

2.2 Empirical Associations

Empirical Associations are heuristic rules. They may determine what actions should be carried out in certain situations, what lines of reasoning can be disregarded and what conclusions may be drawn from the evidence supplied [Rouse80]. This knowledge level is based upon an expert's opinions and experience. For example, other experts may have conflicting views about this empirical knowledge (not usually written in text books) and it may be probabilistic rather than deterministic.

2.3 Causal Models

Causal Models are collections of interrelated rules usually associated with a particular problem, hypothesis or overall diagnostic conclusion. These rules interact in such a way as to solve problems by searching Causal Chains, therefore solutions are determined by inferring cause-effect relationships [Winston84].

2.4 Strategic Knowledge

Strategic knowledge uses strategies to choose optimum search spaces by the use of meta-rules [Sterling80], that guide a search down a particular path (when several paths exist) before considering any other path. This form of knowledge is basically the inference and search guidance mechanism and is sometimes referred to as "knowledge about knowledge".

3 Knowledge Acquisition by Articulation

There are many reported knowledge acquisition articulation approaches and case studies. These ad-hoc approaches have been developed to elicit rules, heuristics and general domain knowledge. Furthermore, they cannot readily identify the expert's detailed reasoning or behaviour.

Knowledge acquisition by articulation relies upon the interviewing skills and thought processes of the knowledge engineer. The willingness and communication skills of the expert also play a major role in the success of this approach. These requirements are often far from satisfied during a typical knowledge engineer/expert consultation, consequently, the effectiveness of articulation is limited but it can be a useful technique for rapid prototyping. As a result, many iterations of knowledge-base construction/analysis may be required before the expert's domain of expertise is captured to an acceptable level of competence.

3.1 Concept Sorting

Concept Sorting [Chi81] involves extracting a group of concepts from journals, books or glossaries. Each concept is written on an individual card which is then categorised by an expert who describes the relationships between them. Concept sorting simplifies the task of knowledge organisation and forms the basis for producing a hierarchical conceptualization of the domain structure. The success of this technique relies upon the completeness of the concepts originally discovered and written on the cards. There is a high probability that some concepts will be missing or irrelevant initially, since this is done by the knowledge engineer.

3.2. Critical Incident

The Critical Incident memory probing technique requires the expert to focus upon particularly memorable events that require a great deal of expertise [Wellbank83]. Because the selected cases are unusual, it is likely that they will be easy to recall and they have the added benefit of stimulating the interest of the expert. The selected cases are usually extremely difficult or remarkable, therefore it is possible to elicit hidden heuristics and complex rule interrelationships.

This approach relies upon the accurate recollection of the expert's thoughts and behaviour at the time of the event. In practice, it has been found that most experts digress, over-generalise and confuse issues with irrelevant statements and facts. Hence, care must be taken in the analysis of the acquired knowledge, also other elicitation techniques must be used to clarify poorly explained and ambiguous areas.

3.3. Conversation Theory and Focussed Interviewing

Conversation Theory [Johnson85] and Focussed Interviewing [Breuker84] are two elicitation techniques with many similarities. They both resemble normal conversation in which the knowledge engineer probes the domain expert with pre-determined topics and concepts. The effectiveness of this approach depends upon the personality and probing techniques employed by the knowledge engineer who should ideally be proficient in interview protocols so that the expert feels relaxed and unthreatened.

Verbal probes and body language are an integral part of interviewing, for instance a simple "yes" or nod of the head will encourage the expert to continue along the current train of thought. Conversely, if the interviewer shakes his head or says "sorry I don't understand", then this will invoke another explanation or it will direct the expert along another route.

Conversation Theory explores the expert's domain at two levels: level 0 and level 1. Level 0 analysis is aimed at eliciting general knowledge such as descriptions of procedures and algorithms, whereas level 1 analysis is at a deeper level (mainly causal relationships) with emphasis upon explanations of why level 0 procedures work. This approach requires that the knowledge engineer employs "teachback", by the use of verbal probes and body-language, to check his understanding of what has been said. As a result, interviews should be kept short (less than one hour) as constant interruption and repetitive questioning tends to disrupt the expert's train of thought and can inhibit his willingness to co-operate.

Conversation Theory and Focussed Interviewing are similar in many ways, the major features of the latter being that it requires the interviewer to construct a list of topics for discussion so that they can be sequentially investigated. Probing enables the interview to be focussed along a particular thought process (similar to Conversation Theory's two level analysis). The list of topics can be either investigated in a breadth first or depth first approach. A depth first search allows the knowledge engineer to gain an overall high-level understanding of the domain. This is contrary to the natural thought processes applicable to natural conversation in which focussing occurs in an associative depth first manner. Accordingly, a depth first focussed interview is suited to natural conversation and generally provides better results than the breadth first approach.

3.4. Problem Solving

Problem Solving requires that the expert is given a problem which he attempts to solve [Gammack84]. The expert describes the "what" and the "why" (for each goal/subgoal) until the problem is solved. This approach requires that the expert fully understands the reasoning at each step as well as being able to relay his reasoning to the knowledge engineer. This occasionally leads to explanations that are incongruous with general reasoning given by the same expert when confronted by similar problems. As a result, numerous problems must be examined in order to reduce the probability of knowledge misrepresentation.

3.5. Repertory Grids

The repertory grid approach requires the expert to identify a set of domain objects. Relationships between each object are identified by exhaustively selecting the objects in sets of three. The expert then explains how two of the set of three are similar and how they differ from the third. All permutations of the three are analysed and all permutations of objects are grouped and investigated.

This approach is useful for gaining an understanding of the relationships between objects in a specific domain. Unfortunately the repertory grid's Complexity (C) is greater than exponential in the number of objects (the relationship is shown in equation 1). Consequently, repertory grids become unmanageable for domains of more than approximately 10 objects and therefore it has a limited use and is an impractical approach for the majority of domains.

$$C = \frac{n!}{2 \times (n - 3)} \quad \text{-----}(1)$$

where n is the number of objects
in the set and $n > 3$.

4 Knowledge Acquisition by Induction

When a knowledge engineer participates in the process of knowledge acquisition he is basically serving an apprenticeship. In practice, apprenticeships generally require a human to spend a number of years learning by experience and structured courses. The structured courses give the apprentice the background theory of a domain, whereas the experience gained by solving numerous problems eventually leads to the apprentice becoming an expert.

Another example of learning by problem solving is that of undergraduates studying for a degree. Lectures give the student the background theory and any necessary formulas, whereas case studies, laboratory sessions and tutorials teach the student how to apply his knowledge. Only when the student has solved enough problems will he/she be able to pass the relevant examinations. It follows therefore that if case studies and tutorials are used by a knowledge elicitation methodology, then it may be possible to alleviate some of the difficulties associated with knowledge acquisition.

4.1 The ID3 Algorithm

Quinlan's ID3 algorithm [Quinlan83] has stimulated research and commercial development of inductive knowledge acquisition methodologies. This has resulted in packages such as Expert-Ease [Slocombe86], PRISM [Cendrowska87] and NEXTRA [Rappaport88].

This inductive approach requires the expert to initially identify a set of attributes $A = \{a_1, a_2, \dots, a_n\}$. Each attribute can have many associated values $V = \{v_1, v_2, \dots, v_m\}$ that are defined by the expert (such as good, bad and reasonable). An instance of the attribute set A consists of a unique permutation U of the attribute values. For each U, a conclusion or hypothesis d is selected and inserted by the expert. Once the expert has finished inserting all of the possible instance/hypothesis test cases the rules are induced under the following assumptions:

- * all hypotheses are mutually exclusive
- * all attribute values are discrete
- * each attribute is relevant to the set A in order to generate a hypothesis d
- * all hypotheses are correct

The above assumptions and the employed inductive algorithmic approach requires that the training set (test case set) is complete. That is, every permutation of attribute value pairs have been analysed by the induction algorithm. If this is not accomplished

then a degradation in rule set confidence will result. This requirement restricts the ID3 approach to that of small sets of A and small numbers of attribute values in which the Size (S) of the complete training set is defined by equation 2. For example, if A consists of n attributes (n=10), 5 of which have 3 values and 5 of which have 4 values, then S=248,832. This set size is far too large for an expert to analyse due to the complexity of the task in which n decisions have to be made for each instance of S. The Complexity (C) of the demands imposed upon the expert are shown in equation 3, from which the complexity relationship graph (shown in figure 1) is derived. This graph highlights the impractical demands imposed upon the expert's analytical skills, the exception being in domains such as medical diagnostics in which large numbers of case studies have been archived.

$$S = \prod_{i=1}^n v_i \quad \text{--- (2)}$$

$$C = K \times n \times \prod_{i=1}^n v_i \quad \text{--- (3)}$$

where K is a constant.

The ID3 rule induction algorithm uses entropy functions [Shannon49] to develop a rule decision tree. An entropy function is a measure of the information yielded by an event or as a measure of how unexpected the event was. In this approach entropies depend only upon the probabilities of the events considered, therefore each attribute value in the set A is considered as information with some bits of information deemed more important than others. Intuitively, it would not be unreasonable to assume that the entropy of an event (H) has the following properties:

$$H \text{ is nonnegative } \forall p \in]0,1] \quad (H(p) \geq 0) \quad \text{---(4)}$$

consequently information about an event does not reduce the information content already known.

Furthermore, H is additive:

$$\forall p, q \in]0,1] \quad (H(pq) = H(p) + h(q)) \quad \text{---(5)}$$

and therefore the information obtained from two independent events is the sum of the information yielded by the two single events.

Another feature of H is that it is normalised so that:

$$H(1/2) = 1 \quad \text{---(6)}$$

this supposition (the simple alternative) implies that if an event has a probability of occurrence equal to 1/2, then its opposite is equally probable.

The above assumptions lead to Shannon's entropy function for complete probability distributions, in which the entropy of a set of classifications is defined as:

$$H = - \sum_i p(d_i) \log_2 p(d_i) \quad \text{----(7)}$$

where $p(d_i)$ is the probability of hypothesis d_i occurring.

ID3 uses entropy to generate a rule tree structure by recursively partitioning the training set of examples into a number of subsets of A . This is done in such a way as to maximise H , whereby $H=0$ in the final recursive step. This systematic approach ensures that each partition selects the attribute that yields the most information, therefore the tree is created in which maximum information yield is achieved by a decision (attribute value pair) defined in the node nearest to the root of the decision tree. Accordingly, the ordered insertion of decisions based upon entropy ensures that a minimal decision tree is formed and a compact set of rules induced.

A typical example of a generated decision tree structure, consisting of attribute value pairs and conclusions is shown in figure 2. From this structure the following rules can be derived:

- 1) $a1, v1 \rightarrow d1,$
- 2) $a1, v2 \wedge a2, v1 \rightarrow d2,$
- 3) $a1, v2 \wedge a2, v2 \rightarrow d1,$
- 4) $a1, v3 \wedge a2, v1 \wedge a3, v1 \rightarrow d2,$
- 5) $a1, v3 \wedge a2, v1 \wedge a3, v2 \rightarrow d3,$
- 6) $a1, v3 \wedge a2, v2 \rightarrow d1.$




Figure 1. Complexity Relationships of Rule Induction Training Sets




Figure 2. Example of a Decision Tree Structure

4.2. The Disadvantages of the ID3 Algorithm

The ID3 decision tree rule induction approach is a computationally efficient tool for deriving rule-based systems. Unfortunately, it has a number of undesirable features. Firstly, the quality of the derived rules depends upon the precision and completeness of the training set of test cases. Secondly, decision trees are inherently difficult to analyse, therefore it may be necessary for an inference engine to search the complete

structure before a hypothesis can be identified. Finally, the major limitation of decision trees is that they rely upon a knowledge-base consisting of a single tree structure. This is often an inadequate model for numerous domains of expertise and as a result, redundant rules have to be incorporated into the tree. Unfortunately, this leads to excessively large knowledge-bases and instigates the generation of redundant questions. For example, consider the rule structure of figure 2, if the ES knows that a2,v2 is true then it is not necessary to analyse any other attribute value pairs because a1 and d1 can be concluded regardless of any further information. With this decision tree example the value of a1 is required before a conclusion is reached which in certain domains can lead to unnecessary, costly and time consuming investigations.

To overcome the limitations of decision tree structures, modifications to the ID3 algorithm have been investigated. One approach, PRISM [Chendrowska87] uses a measure of information gain instead of entropy. The major difference of this modified ID3 algorithm is that PRISM identifies only relevant attribute-value pairs, whereas ID3 will identify the most relevant attribute and utilise each of its values even if one or more of these values are redundant. Although PRISM provides a flexible rule structure, it still relies upon the completeness and accuracy of the training set, therefore rule induction techniques have a limited scope for practical applications.

5 Structured Automated Knowledge Acquisition

Automated knowledge acquisition approaches attempt to eliminate the need for highly skilled knowledge engineers. This is accomplished by constraining and guiding the expert's reasoning processes by interactively processing formatted information offered by the user, so that the automated acquisition tool can store the knowledge in a highly specific representation scheme. This ensures that the acquired knowledge is captured in a format suitable for an expert system shell problem solving mechanism.

The foundation of the automated knowledge acquisition approach is that of identifying the functions of a domain's knowledge and partitioning it into a number of functional modules. This enables the knowledge engineer to guide the articulation process, as opposed to asking unstructured questions which rely upon the co-operation and communication skills of the expert. Another role of knowledge partitioning and knowledge function identification is that of explaining to the user the knowledge that was used in the route to a conclusion. This can be considered as the inverse of knowledge acquisition in which a human imparts expertise to the system, whereas when the system is explaining its reasoning it is imparting knowledge to a human.

After the function of a knowledge module is defined and understood, it is possible to identify any incomplete or missing associated areas. Furthermore, the functionally partitioned modular knowledge-base structure, in which each partitioned block is allocated a task or set of tasks, simplifies the design, debugging and evaluation of the evolving ES.

5.1 MORE

Over the past decade there have been many reported automatic knowledge acquisition tools such as TIERESIAS [Davis82], ETS [Boose84], MORE [Kahn85] and MOLE [Eshelman87]. During an interactive knowledge acquisition session these tools generally construct the knowledge-base in a step by step process. At each stage the acquired knowledge becomes more complete therefore allowing the tool to focus upon relevant questions. Inconsistencies are also highlighted and modified as required, for example, MORE is founded upon a three stage process in which stage one is a simple acquisition of observations and their corresponding hypotheses. Whilst the expert is inserting these relationships MORE will request further information so that the

acquired knowledge will be of a specific nature. At this stage a certainty factor associated with the individual rules may be included.

The second phase of MORE's knowledge acquisition process is the analysis of the knowledge-base that was constructed in phase one. This involves the evaluation of each hypotheses' certainty factor. If a hypothesis occurs only once in the knowledge-base and its certainty factor is neither for or against the hypothesis, then this hypothesis is investigated further by MORE. Hence, meaningless or vague hypotheses are eliminated and redundant rules removed from the knowledge-base.

In the third automated acquisition step, inconsistencies are investigated so that the certainty factors associated with highly specific rules have stronger weights than less specific ones. For example, consider the following two rules:-

If A and B and C then D.

If C then E.

The certainty factor associated with the first rule should be higher than that of the second rule, if not, there has probably been an error in the expert's reasoning during phase one or two. MORE checks for these inconsistencies and reports them to the expert so that modifications can take place.

The designers of these automated acquisition approaches have claimed that they have practically eliminated the "Feigenbaum Bottleneck" [Feigenbaum79] and have dramatically reduced the need for highly experienced knowledge engineers. In practice this is not the case and although these tools are a step in the right direction, further research is required to improve their performance; once this is achieved, it may be possible to use these tools to rapidly acquire highly specific domain dependant knowledge. This will only be possible if a knowledge engineer is involved in the preliminary and final stages of the knowledge acquisition process.

6 Conclusion and future directions

There are many problems associated with knowledge acquisition. For example, articulation approaches are iterative and require highly skilled interviewing techniques so as to maximise the interactions between the knowledge engineer and domain expert.

In contrast to articulation approaches, rule induction techniques are non-iterative but are generally only applicable to small sets of relationships, the exception being when large amounts of case studies are accessible. Furthermore, unless an exhaustive training set is used other acquisition techniques may be required to refine the set of induced rules.

Automatic knowledge acquisition methodologies are becoming more evident in ES construction. Unfortunately, they require the domain of expertise to be modelled by highly specific formations (rule-based structures with a probabilistic measure of belief). Currently automated approaches have limited capabilities and they do not explicitly utilise hierarchical relationships or object oriented structures. Thus, if these features are incorporated into an automatic acquisition tool, driven by "question and answer" interactions, then highly structured knowledge schemas could be created. Such an approach could be greatly facilitated by an integrated software architecture consisting of a new hierarchical knowledge representation schema. This would be implemented by the use of a graphical interface and a constrained natural language processor. As a result, techniques applicable to graphical modelling and application specific natural language processing must be investigated.

7 REFERENCES

[Boose84]

Boose J., *Personal construct theory and the transfer of human expertise*, Proc. Fourth National Conf. on Artificial Intelligence, Austin, Texas, 1984.

[Cendrowska87]

Cendrowska J., *PRISM: An algorithm for inducing modular rules*, Int. J. Man-Machine Studies, Vol 27, pp 349-370, 1987.

[Chi81]

Chi M.T.H., Fetlovic P.J., Glaser E., *Categorisation and Representation of Physics Problems by Experts and Novices*, Cognitive Science, 5, (1981), pp 121-152.

[Davis82]

Davis R., Lenat D., *Knowledge-based Systems in Artificial Intelligence*, McGraw-Hill, New York, 1982.

[Eshelman87]

Eshelman L., McDermot J. *MOLE: A knowledge acquisition tool that uses its head*, Proc. Fifth National Conf. on Artificial Intelligence, Philadelphia, Pennsylvania, 1985.

[Feigenbaum79]

Feigenbaum E. A. *Themes and case studies in knowledge engineering*, In D. Michie (Ed.) *Expert systems in the micro-electronic age*. Edinburgh: Edinburgh University Press, 1979.

[Gammack84]

Gammack J.G., Young R., *Psychological techniques for eliciting expert knowledge*, 5th Technical Conference of the British Computer Society Specialist Group on Expert Systems, (University of Warwick, 1984), pp 105-112.

[Johnson83]

Johnson L., Keravnou E., *The importance of knowledge representation schemes in the performance of expert consultant systems*, MCSG/TR/29, (Dept. of Comp. Sci. Brunel University, 1983).

[Johnson85]

Johnson L., Johnson N.E., *Knowledge elicitation involving teachback interviewing*, (Man Computer Studies Group, Brunel University), 1985.

[Kahn85]

Kahn G., Nowlan S., McDermot J. *MORE: An intelligent knowledge acquisition tool*, Proc. Ninth International Conf. on Artificial Intelligence, Los Angeles, California, 1985.

[Rappaport88]

Rappaport A.T., Gaines B.R., *Integrating knowledge acquisition and performance systems*, AAAI-88 Workshop on Integration of Knowledge Acquisition and performance systems, St Paul, USA, 1988.

[Rouse80]

Rouse W.B., Rouse S.H., Pellegrino S.J., *A rule-based model of human problem solving in fault diagnosis tasks*, IEEE Transactions on Systems, Man, and Cybernetics, Vol 10, No 7, (1980), pp 366-375.

[Shannon49]

Shannon C.E., Weaver W., *The mathematical theory of communication*, University of Illinois Press, 1949.

[Slocombe86]

Slocombe S., Moore K.D., Zelouf M., *Engineering expert systems applications*, Expert Systems '86, Brighton, 1986.

[Sterling80]

Sterling L., *Implementing problem solving strategies using the meta-level*, Dept. of Artificial Intelligence, (University of Edinburgh, 1980).

[Wellbank83]

Wellbank M., *A Review of Knowledge Acquisition Techniques for Expert Systems*, (British Telecom. Research Laboratories, Martlesham Heath, Ipswich), 1983.

[Winston84]

Winston P.H., Prendergast K.A., *The AI Business*, (MIT press, 1984).

THIS IS A BLANK PAGE

DOCUMENT CONTROL DATA SHEET

Page Classification UNCLASSIFIED
Privacy Marking/Caveat N/A

1a. AR Number AR-006-715	1b. Establishment Number ERL-0553-RN	2. Document Date JUNE 1991	3. Task Number DST 90/151	
4. Title TECHNIQUES FOR KNOWLEDGE ACQUISITION		5. Security Classification	6. No. of Pages	19
		<div> <div>U</div> <div>U</div> <div>U</div> </div> Document Title Abstract	7. No. of Refs.	17
		S (Secret) C (Conf) R (Rest) U (Unclass) * For UNCLASSIFIED docs with a secondary distribution LIMITATION, use (L)		
8. Author(s) S. J. Cosgrove		9. Downgrading/Delimiting Instructions N/A		
10a. Corporate Author and Address Electronics Research Laboratory PO Box 1600 SALISBURY SA 5108		11. Officer/Position responsible for Security.....N/A..... Downgrading.....N/A..... Approval for Release.....D/ERL.....		
10b. Task Sponsor				
12. Secondary Release of this Document APPROVED FOR PUBLIC RELEASE Any enquiries outside stated limitations should be referred through DSTIC, Defence Information Services, Department of Defence, Anzac Park West, Canberra, ACT 2600.				
13a. Deliberate Announcement NO LIMITATION				
13b. Casual Announcement (for citation in other documents)				
<div> <input checked="" type="checkbox"/> No Limitation </div> <div> <input type="checkbox"/> Ref. by Author & Doc No only </div>				
14. DEFTEST Descriptors Knowledge Acquisition Expert Systems			15. DISCAT Subject Codes 120901	
16. Abstract Knowledge acquisition is commonly regarded as the major bottle-neck in Expert System (ES) development. Numerous techniques and commercial packages have been built in order to widen this bottle-neck. Unfortunately, these have had little success, except in simple, highly confined domains, therefore knowledge acquisition is still regarded as a "black art". This paper critically analyses current acquisition techniques and suggests future research directions.				

Page Classification
UNCLASSIFIED
Privacy Marking/Caveat
N/A

16. Abstract (CONT.)		
17. Imprint Electronics Research Laboratory PO Box 1600 SALISBURY SA 5108		
18. Document Series and Number ERL-0553-RN	19. Cost Code 820441	20. Type of Report and Period Covered ERL RESEARCH NOTE
21. Computer Programs Used		
22. Establishment File Reference(s) N/A		
23. Additional information (if required)		